Section 2

The Laboratory has been heavily vested in supercomputing since its founding

That tradition continues today

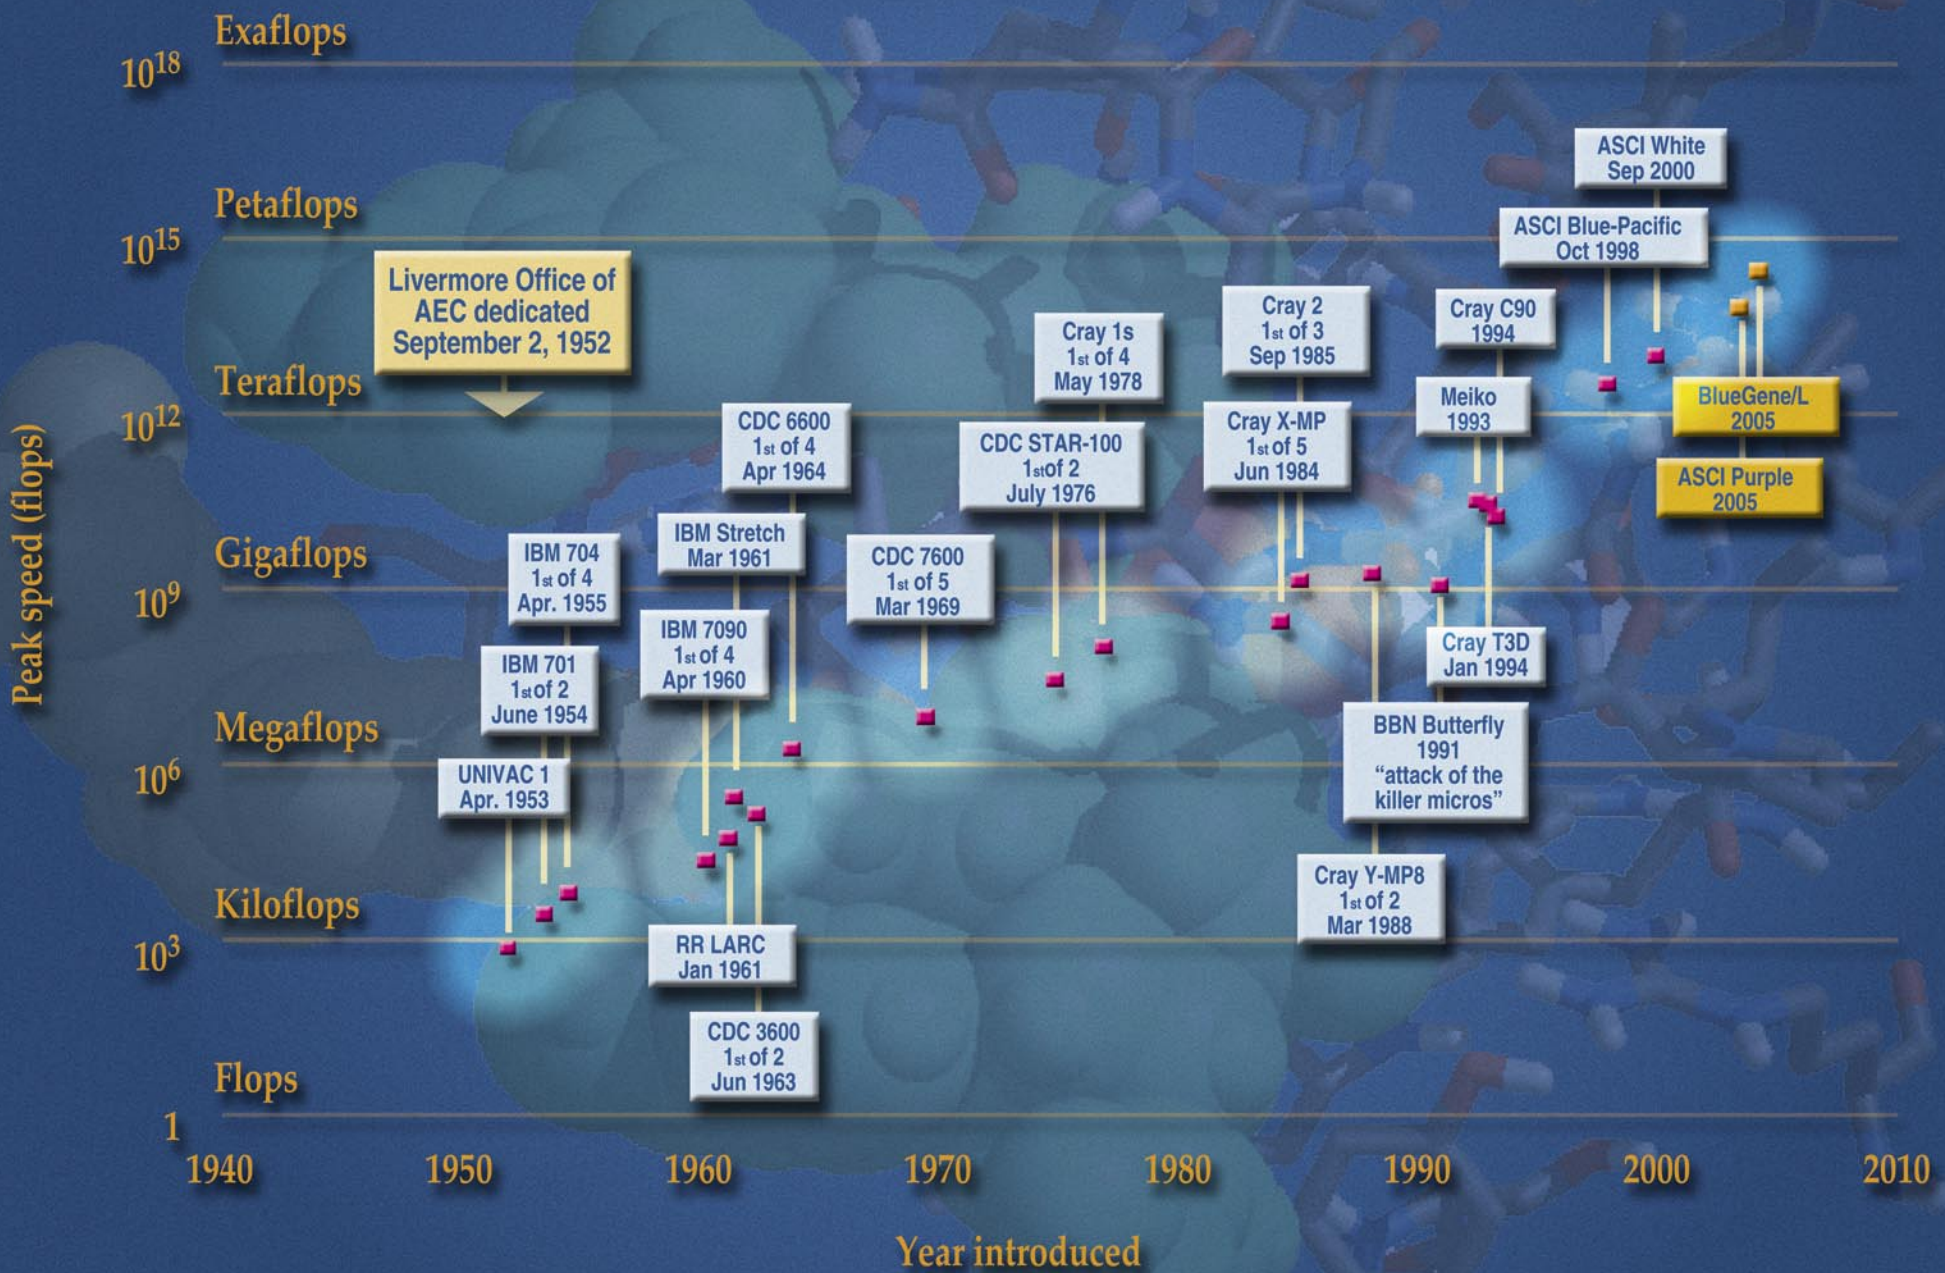

2.00 Providing the Desktops to Teraflops Computing Environment

The Computation Directorate provides LLNL with a world-class computing and networking environment capable of meeting laboratory mission and program needs. We support the Laboratory's computing and communications infrastructure, spanning users' needs from the desktop to the high-performance computing platforms. We assume responsibility for planning and operating the scientific computing facilities, developing tools that enable effective use of these facilities, providing expertise in desktop support, and running as site-wide network backbone for both classified and unclassified systems. We also undertake essential computational, communication, and computer security research required to sustain this computing environment.

At the highest level, two broad objectives are ubiquitous. First, we seek coordination of services so that the connections between the office environment and the HPC infrastructure become increasingly transparent. Second, from a strategic perspective, we require a framework for a petaflop strategy. Such a framework includes both strategic investments in simulation environments as well as in innovative architectures. The body of this Section reflects both of these objectives.

Sustaining a world-class scientific computing environment demands careful balancing of system components and planning of changes to exploit constant improvements in the technology. This environment consists of more than just the latest supercomputers. It includes production computing resources, fast I/O subsystems, high-capacity archival storage facilities, and high-speed network interconnects to link all of these components. Future planning includes detailed predictions of our expected needs and close collaborations with industry to tailor appropriate and cost-effective solutions. The first two reports in this Section describe our accomplishments and strategic planning for building, operating and evolving our high-performance computing systems and networks.

As these computing systems grow in size and complexity, the challenge of using them efficiently and effectively grows as well. To help users address this challenge, we develop tools, system software, and an application infrastructure, again in partnership with industry. We also provide user services, training, documentation, and consulting. The next two reports in this Section describe some of our most critical activities to support our users' ability to use the large systems well. These systems can generate terabytes and petabytes of data far faster than we can view and assess. One of the reports describes our latest development and deployment activities for visualization and data assessment. The next report describes additional support for HPC users, including access to information on how to use these systems, optimization tools and techniques, and strategies for submitting and tracking large production jobs that could take weeks or months to complete.

The last three reports in this Section complete the landscape needed by users. They must be able to work from their desktops, use backbone networks to access the HPC resources and external networks, and do everything in a safe, secure manner. The desktop environment efforts need to anticipate, integrate, communicate, and implement the information technology requirements of LLNL's programs and the institution. This includes technical support for Macintosh, Window, and Unix systems and servers, local help desks, and Web page development services. The network backbone must provide secure, reliable, effective access to information and computing resources from the desktop by delivering networks, centralized system administration services, and centralized enterprise services. Cyber security continues to be extremely challenging because we must maintain our ability to communicate and learn via Internet access and at the same time fend off increasingly sophisticated and persistent efforts to gain access to our internal networks.

Figure 2.00-1. The Multiprogrammatic Capability Resource combines open source software with cluster architecture to provide Advanced Simulation and Computing-level supercomputing power for unclassified research.

2.01 Platform Strategy and Systems, in Production and Planned

Problem Description

Our strategic and industrial collaborations in HPC center on delivering computing platforms to production environments in support of programs of national interest, with an eye to enabling realistic ramps to petaflop-scale systems. To that end, we work with multiple sources of computing technology to judge the boundary between promises and real computing capability, as well as to distinguish industrial trends. Based on these interactions with the computing industry we have developed a straddle strategy to deliver technology to the ASCI program and to the institution.

Technical Approach/Status

The strategy is depicted in Figure 2.01-1. The ASCI program has enjoyed success by riding Curve #1, the **Proprietary, Vendor Integrated SMP** cluster technology. As part of the ASCI tri-laboratory complex, LLNL and LANL played a significant role in establishing this technology path with the ASCI Blue procurement. This approach has taken the tri-lab HPC community from about 50 GF in CY1995, to 12.3 TF on ASCI White in late CY2000, and all the way to 20 TF on ASCI Q in late 2003. Curve #1 price–performance, however, is being eclipsed by Curve #2, **Open Source Commodity Clusters** with the Linux operating system (Beowulf technology).

In addition, on Curve #3, **Innovative Concepts**, such examples as IBM's system-on-a-chip (SOC) technology for embedded applications also have the potential for extreme price performance. Each of these curves has a different price–performance and risk trajectory. The SOC design shows high risk for ASCI-scale platforms today, but holds the promise of an affordable petaflop compute engine in the 2006–2008 timeframe. Thus, to simultaneously maximize the benefit and optimize the potential of emerging technologies, while also extracting benefits from mature technologies, we are engaging development of platforms and software on all three curves. When this strategy was launched, the intent was to move the most important programmatic work onto Curves #2 and #3 only when these technologies had matured, featured reduced risk, and were well understood as production environments.

During 2003, we saw a positive outcome of this approach. Experience with Thunder (IA64) and MCR (IA32) showed LLNL that Curve #2 could potentially deliver at ASCI scale with medium risk. We indicate this change in the trajectory by moving the curve to the left from the dashed blue (old) to the solid blue (new). With this new price–performance reality in mind, the Laboratory has renegotiated the ASCI Purple contract with IBM. This approach represents a win–win, because it accelerates IBM's trajectory in fielding very large, low-cost systems with the Power series, expanding the space of such solutions beyond INTEL and AMD. This will help maintain healthy competition and low costs.

Progress in 2003

A number of powerful systems were sited at LLNL in 2003—Thunder, Lilac, Violet, and Magenta. Other systems have been brought into production, MCR and ALC; and others have seen major upgrades — the Penguins, Adelie and Emperor. Much of this has been accomplished with the CHAOS (Clustered High-Availability Operating System) software stack.

In short, CHAOS augments the standard Linux Red Hat distribution with support for HPC clusters, including scalable system management and monitoring tools (primarily developed at LLNL), a high-performance interconnect (Quadrics Èlan3 and Èlan4), the Lustre parallel file system from CFS (heavily funded by ASCI) and an advanced resource manage-

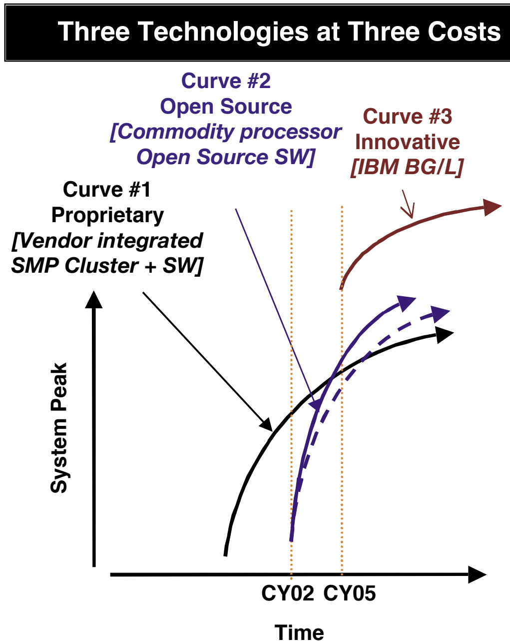

Figure 2.01-1. Price–performance straddle strategy to deliver ASCI technology.

ment and control apparatus (Simple Linux Utility Resource Management or SLURM), also developed at LLNL. These represent the tools necessary to field production HPC clusters, allow LC to integrate new systems rapidly, leverage in-house expertise to provide fast turnaround on bug reports and feature enhancements, and provide a framework for release management. The LC carefully keeps kernel modifications to a minimum to reduce friction with new releases, and leverages its relationship with Red Hat to get new releases into their distribution. It was the experience with this tool on computers of scale that accelerated LLNL's transition onto Curve #2.

Figure 2.01-2 summarizes recent progress. Systems noted in red are either new or have changed significantly in status during the past year, bringing the total peak across all systems on site or under integration close to 90 TF. By this relatively crude metric, LC is currently one of the larger HPC sites in the world.

Looking to 2005, the two IBM contract systems, the 100-TF Purple system and the 180–360-TF BlueGene/L (BG/L) system, seem remarkably well-aligned with the two foci of the new strategy for Advanced Simulation and Computing (still known as ASCI). These foci are: integration into the broader program through providing essential support for stockpile stewardship deliverables, in particular for Directed Stockpile Work (DSW) to support re-certification of weapons systems; and continual reduction in the phenomenology in the weapon simulation codes, including a deeper understanding, in quantitative terms, of their limitations. The latter is particularly important as weapons wander from their test base through aging. From this perspective, the Purple system will provide the must-have, time-critical cycles to the classified program in a highly reliable production environment.

BG/L will provide a research platform to increase prediction by understanding materials properties well enough to reduce phenomenology in the ASCI applications codes. It will contribute at all length and time scales for multiscale materials models (Figure 2.01-3). One can therefore look at the two systems as computational components, each vectored at a different aspect of the program strategy. This approach thus provides a vendor-integrated solution at Beowulf cluster cost and at scales heretofore unattained.

System	Program	Manufacturer & Model	Operating System	Interconnect	Nodes	CPUs	Memory (GB)	Peak GFLOP/s
Unclassified Network								47,847
Thunder	M&IC	California Digital	CHAOS 2.0	Elan4	1024	4,096	8,192	22,938
ALC	ASCI	IBM xSeries	CHAOS 1.2	Elan3	960	1,920	3,840	9,216
MCR	M&IC	Linux NetworX	CHAOS 1.2	Elan3	1152	2,304	4,608	11,059
Frost	ASCI	IBM SP	AIX 5.1	Colony DS	68	1,088	1,088	1,632
Blue	ASCI	IBM SP	AIX 5.1	TB3	264	1,056	396	701
TC2K	M&IC	Compaq SC ES40	Tru64 5.1b	Elan3	128	512	280	683
iLX	M&IC	RAND Federal	CHAOS 1.2	N/A	67	134	268	678
GPS	M&IC	Compaq GS320/ES45	Tru64 5.1b	N/A	49	160	344	277
PVC	VIEWS	Acme Micro	CHAOS 1.2	Elan3	64	128	128	614
Riptide	VIEWS	SGI Onyx2	Irix 6.5.13f	8 IR2 Pipes	1	48	37	24
Qbert	M&IC	Digital 8400	Tru64 5.1b	MC 1.5	2	20	24	25
Classified Network								41,171
Violet (pEDTV)	ASCI	IBM P655	AIX	Federation	128	1,024	2,048	6,144
Magenta (pEDTV)	ASCI	IBM p655	AIX	Federation	128	1,024	2,048	6,144
Lilac (xEDTV)	ASCI	IBM xSeries	CHAOS 1.2	Elan3	768	1,536	3,072	9,186
White	ASCI	IBM SP	AIX 5.1	Colony DS	512	8,192	8,192	12,288
Ice	ASCI	IBM SP	AIX 5.2	Colony DS	28	448	448	672
Blue-Pacific (S)	ASCI	IBM SP	AIX 5.3	TB3	488	1,952	1,164	1,296
Adelie	ASCI	Linux NetworX	CHAOS 1.2	Elan3	128	256	512	1,434
Emperor	ASCI	Linux NetworX	CHAOS 1.2	Elan3	128	256	512	1,434
Ace	ASCI	Rackable Systems	CHAOS 1.2	N/A	128	256	512	1,434
SC Cluster	ASCI	Compaq ES40/ES45	Tru64 5.1b	N/A	40	160	384	235
ICF Cluster	ICF	Compaq ES40/DS10L	Tru64 5.1b	N/A	12	36	12	48
GVIZ	VIEWS	Rackable Systems	CHAOS 1.2	Elan3	64	128	256	717
Whitecap	VIEWS	SGI Onyx3800	IRIX 6.5.13F	4 IR3 Pipes	1	96	96	77
Tidalwave	VIEWS	SGI Onyx2	Irix 6.5.13f	16 IR2 Pipes	1	64	24	38
Edgewater	VIEWS	SGI Onyx2	Irix 6.5.13f	10 IR2 Pipes	1	40	18	24

Figure 2.01-2. Progress summary: LC Systems either in production in March 2004, or currently being integrated, show rapid growth due to the maturation of cluster technologies.

The external community, including scientists from Office of Science and NSF laboratories, reviewed BG/L partnership progress twice in 2003. Both reviews endorsed the rate and quality of progress and encouraged increased engagement with the broad community. The July review panel, chaired by Mike Levine of Pittsburgh Supercomputer Center, reported that this work represented, "an important opportunity to substantially advance both the ASCI mission agenda and the development of very large-scale machines." This computer represents the most complex SOC design that IBM has ever built, yet no critical problems were found with first hardware. The first compute ASICs (application-specific integrated circuits) were delivered June 6, 2003, were booted, and were running applications two weeks thereafter. A 512-node prototype (128th scale of full system, Figure 2.01-4) ran Linpack at 1.412 TF using a slow clock (500 Mhz) rather than the final clock (666–700 Mhz) processors. The MPI latency is expected to come in at 4.3 µs (besting our 7 µs target). Further, IBM is achieving exceptionally high network bandwidth (~80% of theoretical max), accessible even at very small packet size (half the max speed at 500 bytes). Recently, 700-Mhz processors have been run using 512 MB of memory, increasing probability that the machine can be used for a wide variety of applications.

Early in 2004, Sandia, LANL and LLNL, with the concurrence of DOE HQ, agreed to move forward with the parts order in 2004. All three laboratories are targeting science applications at the system, to be delivered in phases in the first half of 2005. LLNL is funding the development of five applications, from ab initio molecular dynamics, to crack propagation, to turbulence. ASCI Alliances have been invited to submit proposals for access. LLNL is working with HQ to permit access by Office of Science laboratories collaborating with ASCI.

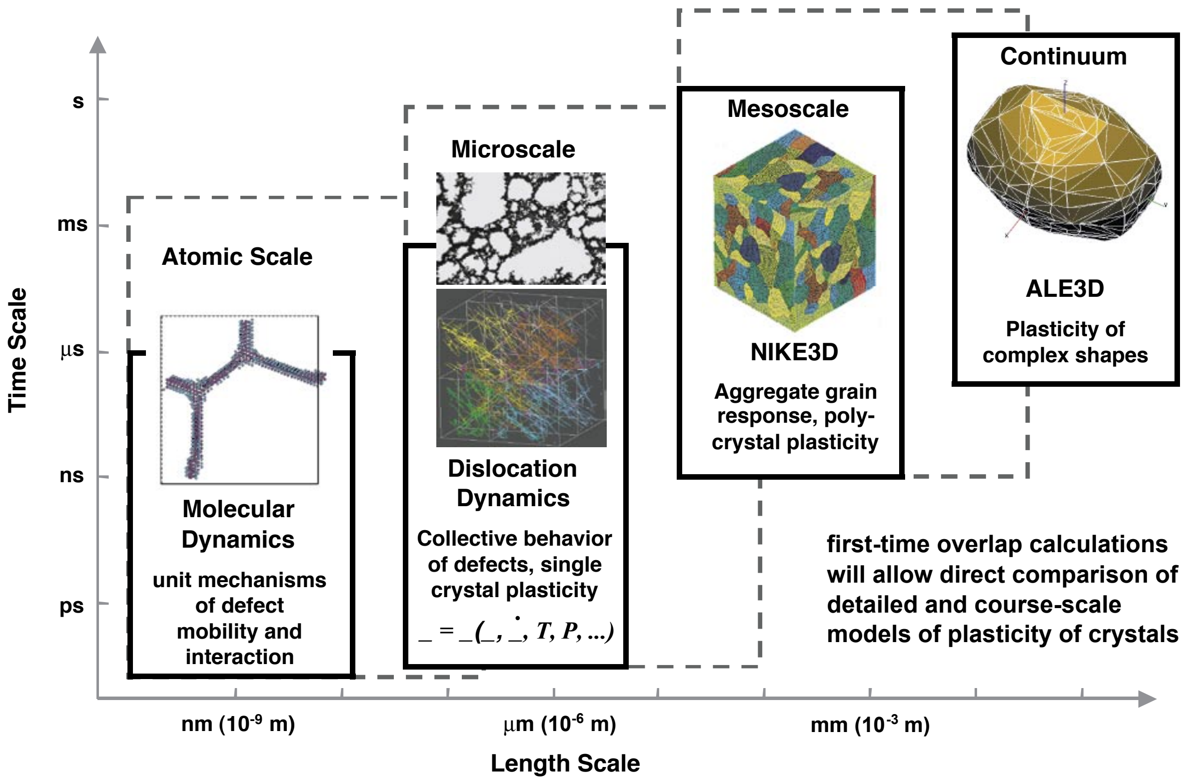

Figure 2.01-3. Multiscale comparisons: Length scales plotted as a function of time.

Figure 2.01-4. 512-node prototype system, operating at IBM Watson Research Center.

Problem Description

Translating LLNL's substantial investments in platforms and applications into successful science requires a balanced computer infrastructure. To ensure this, LC develops an annual **I/O Blueprint**, a planning document that collects platform I/O capabilities plus user requirements, and then presents architecture options, issues, action plans, deliverables, and budget scenarios along with a scope of work for the I/O infrastructure. I/O Blueprints have been used for many years at the LC to assure that investments across all ASCI infrastructure budgets are coordinated. ASCI White success depended heavily on the FY99 and the FY00 Blueprints.

The FY03 I/O Blueprint began by detailing a vision for the next three to five years, centered on an architecture in which computational and visualization resources share a high-performance parallel global file system to provide users with fast, cost-effective uniform access to a very large pool of online storage. This file system will be known as the Site-Wide Global File System (SWGFS), and is shown in Figure 2.02-1. Each of the I/O infrastructure teams began working toward this vision during the year. As outlined below, much progress has been made.

Progress on High-Level I/O Blueprint Deliveries in 2003

In the **networking** arena, LC continues to provide three independent networks to satisfy Labwide I/O needs: a small packet network tailored for interactive access from user desktops, a high-performance four-stripe parallel "jumbo" packet network to facilitate movement of large data sets, and a network for Center-wide NFS access. Anticipating the urgent need for more bandwidth, LC introduced a few 10-Gb Ethernet trunks into production and procured 50 more 10-Gb Ethernet ports. Requirements for greater connectivity were met by deploying newly available high port-density line cards, resulting in over 1300 Gigabit Ethernet ports on the unclassified network. The Visual Interactive Environment for Weapons Simulations (VIEWS) digital delivery effort collaborated with industry to demonstrate a technology for accessing LC's visualization resources from the desktop over the existing networks.

In the **archival** storage arena, LC deployed a new generation of high-performance archive-mover platforms and a Storage Area Network (SAN) disk cache. Upgrading to StorageTek's™ latest generation tape drives tripled tape performance and capacity. The High-Performance Storage System (HPSS) archival software underwent a significant upgrade and an ASCI PSE white paper addressed SWGFS/archive integration. The results of Blueprint-driven changes can be seen in Figure 2.02-2.

Network File System (NFS) upgrades were significant both in the amount of storage capacity provided and in the bandwidth offered. The upgrade plan included the use of RainStorage™ devices, which

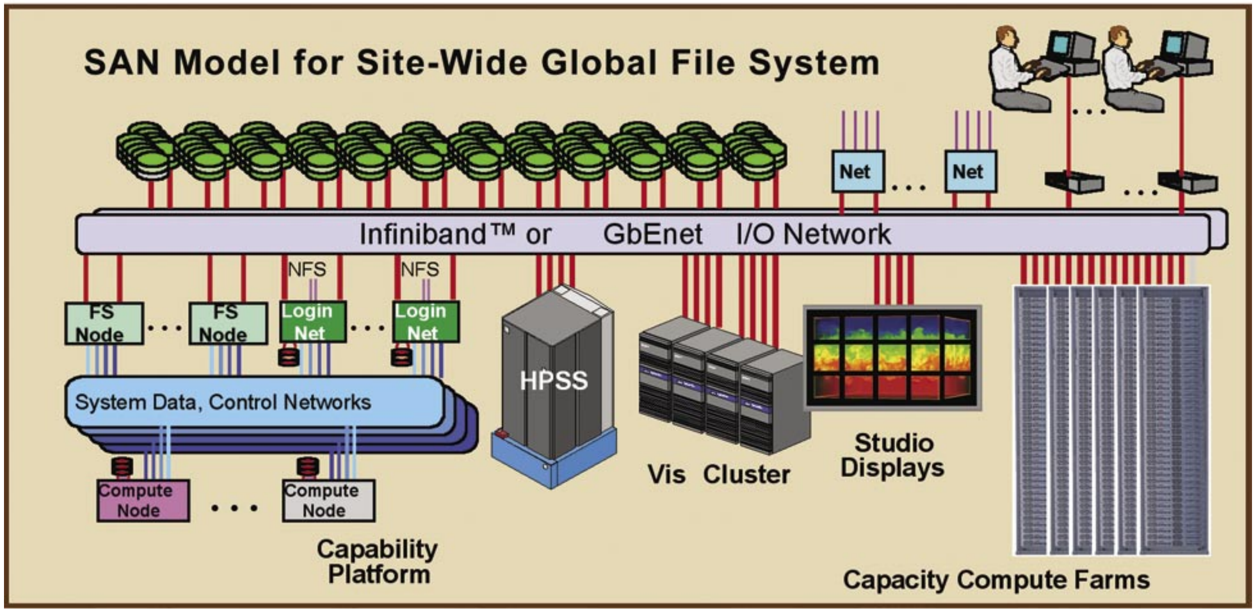

Figure 2.02-1. SWGFS architecture shares one file system, many computers and services.

made it possible to migrate to new NFS servers without unduly impacting normal operations. The upgrade allowed LC to respond to requirements for improvement in multiple areas: the addition of customer scratch space; a significant expansion in home directory space resulting in a large increase in customer quotas; the development of monitoring and planning tools; the deployment of a test platform for new server evaluation; and allowing selected user access to these NFS servers directly from their desktops.

The Center took its first steps toward **SWGFS in production** with the M&IC-funded compute resource MCR, closely followed by the ASCI-funded ALC cluster and the PVC visualization cluster. The Lustre file system, employed by "friendly users" since Fall 2002, is now providing acceptable functionality and performance on these systems. Today, Parallel Visualization Cluster (PVC) and MCR share a single global parallel file system, and the plan is to merge the IP-based Lustre storage infrastructure with the current high-performance parallel network infrastructure and make *all* the storage available to *all* hosts supporting the Lustre file system. This will allow the LC to move toward the visions of "one file system, many computers," and will cut costs significantly.

All of these accomplishments depended heavily on the existence of a flexible test bed infrastructure. The I/O Test Bed proof-of-concept environment continued as a vital tool used by all infrastructure components in identifying and developing reliable high-performance hardware and software solutions. Working together, and focused by the I/O Blueprint, networking, archive, NFS and SWGFS teams were able to provide LC customers with a balanced and world-class production environment for simulation science in 2003. We are also moving forward together toward state-of-the-art services for ASCI Purple and BG/L machines in 2005.

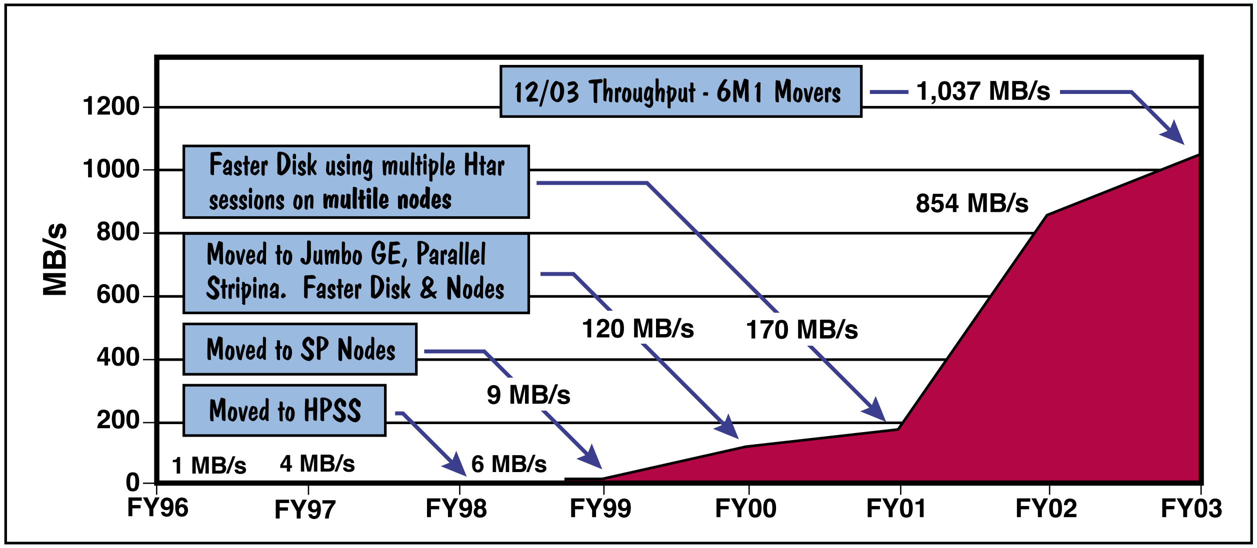

Figure 2.02-2. Maintaining SCF capability platform-to-HPSS performance ratios is an I/O Blueprint requirement.

Problem Description

The very large ASCI computers are prodigious data-generation engines. The I/O infrastructure is the communications and storage component necessary to preserve the data. The final component of the HPC architecture is the data assessment environment.

In support of Stockpile Stewardship, ASCI simulations require hardware and software tools to find, access, manipulate, and visualize the multi-terabyte scientific datasets resulting from large simulations, to compare results across simulations, and to compare between simulations and experiments. Traditional tools cannot cope with the size, scale, and complexity of terascale datasets. The challenge is to research, develop, and deploy tools that provide users the capability to "see and understand" their data. Although management and visualization of massive datasets is a problem addressed in other scientific and experimental contexts (e.g., satellite images, high-energy physics), accurate analysis of mesh-based ASCI simulation datasets of large magnitude is an ongoing challenge.

The LLNL strategy coordinates ASCI-supported research, development, engineering, deployment and applications support in visualization, data management and data exploration. A major direction targets research and development to create innovative technologies for scientific collaboration, data exploration, visualization, and understanding. Once the requisite technologies exist, they are integrated, tested, and evaluated by a representative set of users. Finally, the technologies are deployed in a generally available, operational and reliable environment for day-to-day use by ASCI users and applications. (Research achievements can be found in Section 4.) Here we concentrate on some of the more notable development and deployment activities for visualization and data assessment, with an emphasis on accomplishments in 2003.

Development and Deployment Activities in 2003

As noted earlier, the transition to a new architecture targeted toward clusters was completed through release of a new software stack. This layered approach provides application toolkits, as well as interfaces to standardized scaling, rendering, compositing, and image delivery libraries, in addition to job and session infrastructure tools. Included were new releases of (DMX) Distributed Multi-headed X11, Chromium, MIDAS and Telepath. DMX is an aggregate X11 server system. Chromium is a distributed, parallel OpenGL application-programming interface based on a dynamically filtered, streaming graphics model. MIDAS is an Open Source tool providing transparent, asynchronous transmission of application-generated imagery from remote visualization servers to desktops. Telepath supports the orchestration of a visualization session, including resource allocation, video switching and delivery and configuration of services.

Together, these packages represent a dynamic visualization applications environment capable of scaling with dataset size, display size and desired levels of performance, using commodity graphics enabled clusters. This software stack was first deployed on the new PVC visualization cluster, LLNL's first production, commodity, PC-based visualization engine. PVC provides direct visualization services for MCR and has already demonstrated its ability to handle multi-terabyte sets from several important codes. The PVC/MCR systems model will serve as a blueprint for future ASCI Purple-related deployments.

During 2003, the TeraScale Browser released its first production, out-of-core, surface-rendering engine, making it possible to render even the largest datasets at the end-user desktop. This release also included the first explicit support of DirectX 9-class graphics hardware in an end-user application, and has been measured to be up to five times faster than the previous release. VisIt, a powerful visualization tool used by DNT, had several phased releases this year, adding improved performance and new capabilities for scalable rendering, stereo, Adaptive Mesh Refinement (AMR) support and movie tools. VisIt was also ported to the ASCI Q machine at LANL, and to Mac OS X this year.

In addition, there was significant progress in the deployment of research tools for data discovery and data query (also see Section 4). Significant work continued to develop and deploy production-quality metadata and directory tool capabilities and fund development of interoperable data models and formats used by large LLNL ASCI code efforts.

Major enhancements were made to SimTracker (a data management tool) for high-level task coordination, remote data management and cross-site data sharing. SimTracker has been deployed with 18 simulation codes. Work has begun to add a comparison framework, change–audit features, and enhanced regression testing to SimTracker. Hopper, a graphical interface supporting HTAR, HPSS, SSH, and FTP, was released in beta form in 2003. It allows compact or detailed views of files and directories, a history mechanism, password management, and a search by name feature. Planned for Hopper in 2004 are content-based metadata searches, and closer integration with LLNL persistent-file transfer tools.

The Center also continued to upgrade ASCI data display capabilities (Figure 2.03-1) with new high-quality projectors and screen in the B132 Data Assessment Theater, a final layout design for the new Terascale Simulation Facility (TSF) Advanced Simulation Laboratory, procurement of high-bandwidth modems for high-resolution user office desktop displays, and evaluation of stereographic projectors.

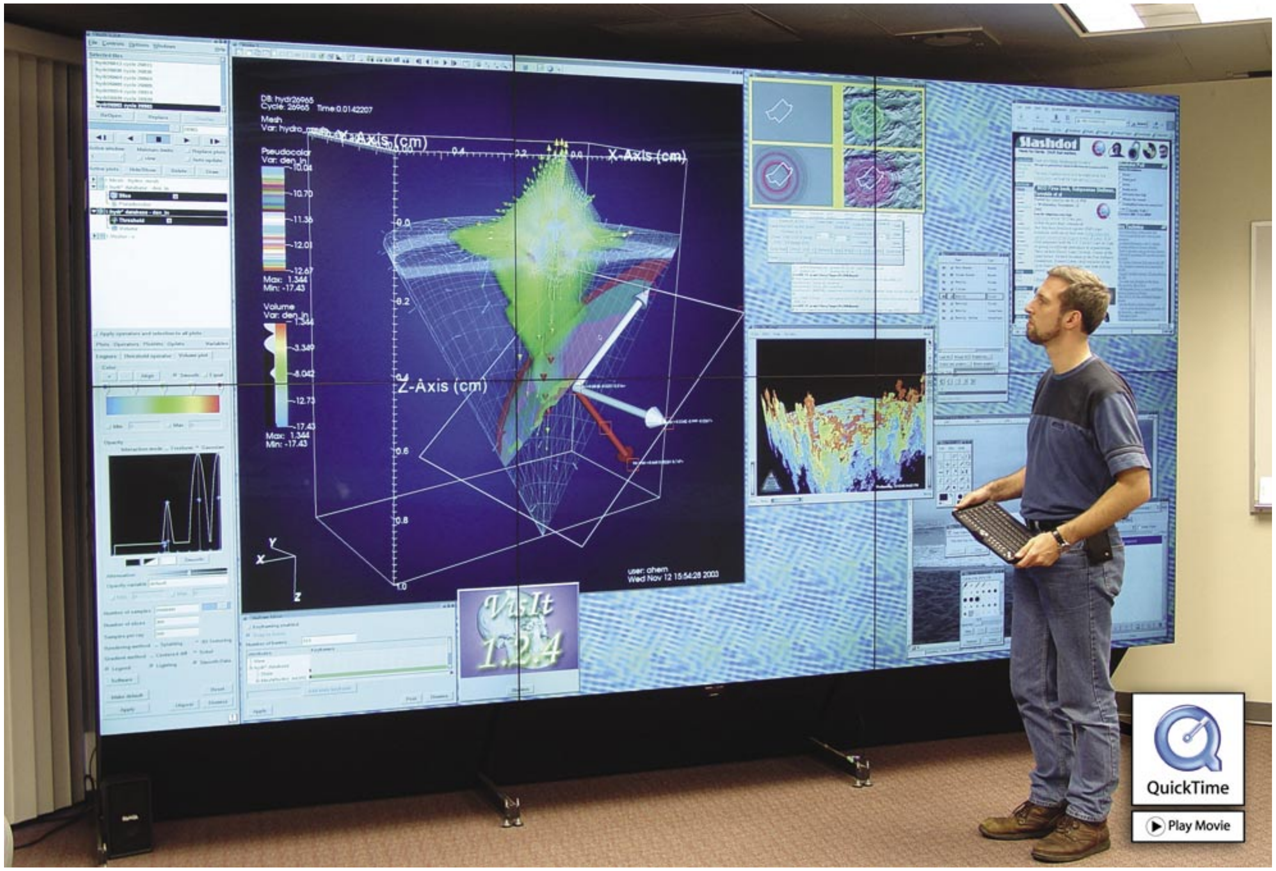

Figure 2.03-1. Interactive data analysis using multiple tools on a cluster-driven display wall.

Problem Description

Success in scientific computing is unlikely in the absence of first-rate consulting and services. Peak speed provides a potential for excellence, but quality service must be commensurate with the hardware on the floor. Our services encompass an "end-to-end cycle of user simulation science." Here, we allude to the cycle of user access, code development, job submission, and data assessment required by computational scientists. We develop tools, system software, and the application infrastructure that enable efficient and effective use of LLNL computing resources. We also provide the user services, training, documentation, and consulting that support the usability, accessibility, and reliable operation of LLNL's computing resources. The stages of this end-to-end cycle are described below.

Technical Approach and Progress in 2003

Stage 1: Enable unencumbered customer access to information to facilitate use of the HPC resources. We provide account services, online documentation, education and training, software quality assurance, and a tracking system for all reported questions and problems. In addition, a high-performance computing (HPC) helpdesk supports the HPC systems, and an institutional helpdesk provides desktop support services and consultation. The Remedy™ tracking system is being upgraded and deployed for several directorates. Both helpdesks ultimately will share this trouble-ticket database. Two such on-line documentation accomplishments for 2003 are noted below.

- Developed software to replace the DynaWeb™. legacy, proprietary document delivery system. The new web-based documentation system was written in XML, is searchable, and currently contains 36 user documents (4,000-plus pages of documentation).
- Replaced the legacy, proprietary document delivery system on the unclassified network with an XML web-based searchable system containing more than 4,000 pages of user documentation, including EZ manuals.

Stage 2: Provide support for scientists porting and optimizing their computer codes. We provide a stable leading-edge parallel application development environment that significantly increases application developers' productivity. Through collaborations with vendors and third-party developers, we ensure a robust environment with the most advanced development tools. We also assist scientists in the code development and maintenance process—compilers, libraries, debuggers, performance measurement and analysis tools, and I/O. Helping customers use these tools effectively speeds up the scientific applications, and therefore optimizes use of machine resources. Specific 2003 accomplishments focused on coding and performance tools, and

- Broadened our code development environment, including compilers, debuggers, performance analysis tools, I/O libraries, and expert consulting services to Linux environments. This contributed greatly to the success of the MCR, ALC, and Lilac Linux/Pentium/Quadrics systems.
- Provided performance tools support, usage expertise, and code analysis that contributed to the successful completion of ASCI code milestones.

Stage 3: Simplify the submission and tracking of production jobs for users. The Livermore Computing Resource Manager (LCRM) tool provides users with a highly sophisticated fair-share scheduler for job and resource management of production computing assets. In addition, LLNL and Linux NetworX are jointly designing and developing SLURM, an open-source resource manager for Unix clusters. SLURM's primary functions are

- To manage a priority-ordered queue of pending work (typically parallel jobs).
- To allocate jobs exclusive or nonexclusive access to compute nodes.
- To provide a framework for initiating, monitoring, and managing jobs.

SLURM and LCRM form a critical part of the LLNL CHAOS cluster software stack, and together assure delivery of resources to the appropriate customer, with very high resource utilization (frequently exceeding 90%) in a contended environment with heterogeneous workloads. Specific 2003 accomplishments included these resource-management advances.

- LCRM v6.9 released April 2003: support for SLURM, better support for prioritizing jobs based on job-size, and replacement of the underlying NQS system with a new, streamlined TBS system.
- LCRM v6.10 released October 2003: support for visualization scheduling, variable job-sizes, and support for pre-emption on IBM systems. Visualization scheduling allows the user to specify a mix of viz and compute nodes.

Stage 4: Provide effective tools for data assessment (including visualization) of production computing runs. We provide extensive visualization and data-management support. Our experts develop and support tools on a wide range of platforms for users representing many disciplines, and for offices, theaters, work centers, and conference rooms. Services include consulting on scientific visualization packages, data-management tools and graphics utilities; authoring and producing movies and DVDs; demo support for PowerWalls; and the stewardship of three theaters. Specific 2003 accomplishments for data assessment included these.

- Produced visualizations in support of M&IC science runs and ASCI milepost reviews.
- Documented computational science results with Science of Scale video (with new material for SC2003).
- Supported numerous high-level demonstrations at all three PowerWalls for audiences from congressional staffers to military dignitaries.

Customer Profile, 2000 through 2003

Customer Profile	2000	2001	2002	2003
Number of Active Users	1959	2219	2304	2450
Classified	964	1080	1231	1360
Unclassified	1609	1826	1847	1904
Number of Remote Users	576	676	648	709
Sandia	85	128	133	163
LANL	70	88	108	119
ASCI Alliances	116	122	110	120
Other	305	338	297	307
Average Number of Hotline Contacts per day	100	110	116	109
4HELP Number of calls per day	63	65	62	77
Average Number of accesses to web pages per day	1080	7145	7308	8675
WWW Documentation				
Documents Available	33	37	35	37
Number of pages of documentation	5131	3798	4315	4331
Compiled Customer Assessments				
Number of classes offered	13	3	11	6
Class Evaluation	Avg. score 4.5 out of 5. Comments Favorable	4.9 out of 5.	4.7 out of 5.0	4.4 out of 5.0
Number of users registered	374	377 registered, 235 attended	514 registered, 426 attended	300
Number of sessions offered	33	20	34	41

Figure 2.04-1. Customer profile metrics help LC to monitor the quality of its services.

LC's success is evidenced by the customer profile metrics gathered over time (Figure 2.04-1), and the high effective-node utilization of the machines. Figure 2.04-2 shows how quickly our services and software helped users effectively utilize MCR when it opened to them; this high level of utilization continues today.

We also improved institutional user services significantly in 2003. Our institutional helpdesk acquired Open LabNet, One Time Password, and all remote-access support for the Laboratory. We increased phone coverage from two to three staff during the day, to four at all times during the second half of the year. Although call volume increased more than 14%, the abandonment rate dropped from 15.64% to 8.88%. In addition, during the second half of 2003, call wait dropped 38%. Remedy™ was successfully rolled out to NAI, NIF, Computation, and parts of Engineering, allowing all of these directorates to pass trouble tickets and gather important support metrics.

Our goal is to work with all LC customers so that we meet their time-critical requests to their full and complete satisfaction.

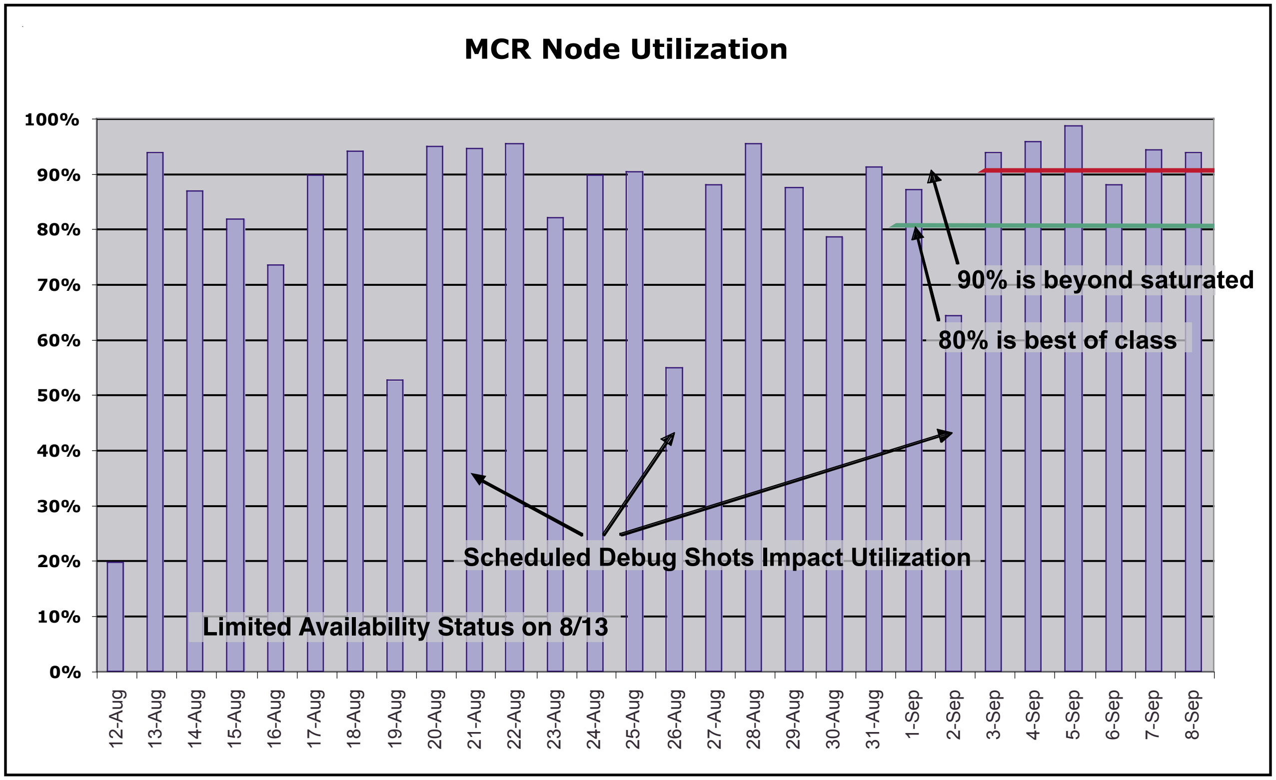

Figure 2.04-2. Complex scheduling algorithms deliver cycles appropriately and guarantee high utilization.

2.05 Desktop Support

Problem Description

All computer users, whether for HPC systems or administrative applications, start at the same place: desktop systems. We provide a competent, trained staff capable of supporting the Lab's wide variety of desktop platforms and operating systems. We recommend strategies for increased cost efficiency or improved performance to our users on an ongoing basis. We also give our users the information they need to make timely, informed decisions about their hardware and software budgets. Significant progress was made on several fronts this year all aimed at increasing desktop computing productivity for LLNL customers across all 13 Laboratory directorates.

Technical Approach/Status

An outside independent benchmarking company surveyed LLNL customers in early 2003. The benchmarking company made several high-level recommendations on ways to improve support. The recommendations were to focus on increasing support responsiveness, support availability and support expertise. An action plan was developed in response to the recommendations and improvements were targeted to areas that had particular issues. In one area, hours of the local support were expanded in direct response to customer feedback. To help solicit satisfaction feedback on an ongoing basis each support area has instituted a 30-second satisfaction survey on closed jobs. The survey is not sent to users on every job completion, but is sent periodically to do spot checks and solicit feedback to uncover issues on a more regular basis. Additionally, expected response-time matrices were created for each support area that indicate an expected problem resolution time by job priority and also define job priority levels. These matrices are intended to create a more consistent customer understanding of priority levels and response times. The final mechanism employed to improve support responsiveness is that technicians in most support areas are calling a handful of customers per week and personally communicating to them when they can expect their problem to receive attention. Our belief is that, in general, customers are more satisfied waiting a week to have their problem resolved when a reasonable cause for delay is communicated to them. Feedback from customers on the personal follow-up has been uniformly positive.

LLNL embarked on a "Total Cost of Ownership of Distributed Computing" study done by an independent benchmarking company. The preliminary report has been received and includes the following recommendations: provide stronger centralized IT Governance; strengthen the centralized help desk and services, implement more standards in the area of desktop operating systems and browsers; and use lifecycle management best practices to drive costs down. Another recommendation is to take a hard look at the number of desktops per customer, to ensure that our 2:1 ratio is required. We have not yet developed an action plan in response to these recommendations, however none of these recommendations comes as a complete surprise. We believe our overall strategy addresses most of them.

Progress in 2003

A major milestone was achieved in our effort to provide an institutional Active Directory (AD) production forest. Active Directory is Microsoft's directory service designed for distributed computing environments. AD allows organizations to centrally manage and share information on network resources and users, while acting as the central authority for network security. In addition to providing comprehensive directory services to a Windows environment, AD is designed to be a consolidation point for isolating, migrating, centrally managing, and reducing the number of directories that companies require.

A forest is the AD entity that is the root for a group of domains and organizational units. It is key to implementing computer security efficiently across LLNL desktops. The forest is up and functioning for two large Directorates (NIF and DNT), and several smaller support areas have migrated large portions of their environments to the institutional forest.

The past investments in our centralized tools have been paying unexpected dividends in the wake of the onslaught of Windows OS security issues. Radia, LLNL's Automated Software Delivery (ASD) tool now resides on more than 7000 LLNL PCs. It is routinely used to provide timely software updates. The NIF Directorate used Radia to deploy the blaster patch to 75% of 1200 targeted PCs in one recent overnight distribution. Even with this success, it is true that ASD is not designed to be an OS patching tool and requires substantial resources to use as a patching solution. Because of this, one of the efforts we have initiated is a search for a cost-effective Windows patching solution.

We have also increased our focus on technology strategy and vendor relationships. The concentrated effort in this area resulted in generation of the "Role of Technology Watch" document (Draft), visits to Apple, Microsoft and Dell to discuss roadmaps and strategies and an ongoing dialog with these vendors to fully understand and optimize the benefits available through existing maintenance contracts. We expect the technology watch role to expand to include focused programmatic requirements gathering, targeting of new technologies to particular advances in program ROI and Pilot studies of promising technologies.

This effort dovetails nicely with our efforts to formalize and publish LLNL recommended software, hardware and browser support strategies. Significant headway was made this year. A proposed strategy to formalize and publish support strategies has been created and vetted with the Desktop Advisory Group (DAG), and various support organizations at LLNL. The strategy includes four phases in the product lifecycle (target, current, containment, and no support) and includes definitions of what "support" means to various service providers at each phase in the lifecycle. The strategy is a multi-year plan for migrating to new software versions and moving off of old versions as vendors drop support for them. Figure 2.05-1 is the draft LLNL PC Operating System strategy. By providing a suggested road map to the institution, we hope to help programs proactively plan their desktop and software purchase and retirement decisions more effectively as well as give application providers a target to shoot for when developing or purchasing institutional applications.

Significance

In summary, much of the year has been spent conducting benchmarks and establishing a context from which to make ROI decisions while simultaneously pursuing improvements in infrastructure and tools (AD, ASD, patching), standards (software, hardware, and browser support strategies) and keeping an eye on new and relevant technologies. We continue to lay a strong foundation for continuing improvement in the out years. Strong central IT governance is fundamental to our ability to substantially increase ROI for our customers. This governance is vital to improving and consolidating the effort involved in effective communication, solicitation, and receipt of buy-in from more than 44 desktop support funding sources and many more stakeholders at LLNL.

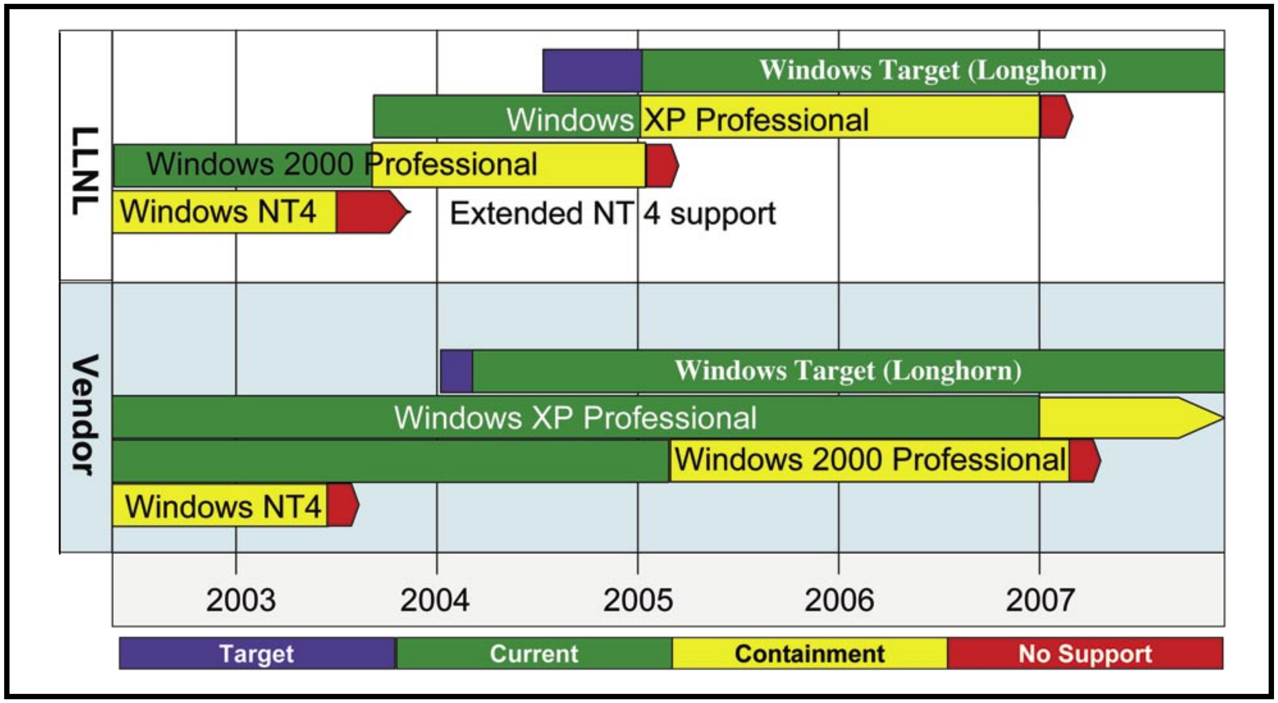

Figure 2.05-1. Windows operating system support lifecycle.

2.06 Laboratory Network Services

Problem Description

Laboratory network services enable effective, legitimate data communications among staff, collaborators, information repositories, and computing resources thus providing a basic set of user services to enhance the efficiency and productivity of all users and staff. A significant challenge is to provide these services ubiquitously and cost-effectively, with an appropriately high level of reliability, performance, and security.

Networks are the foundation of effective, secure communications with and within the Laboratory. Backbone and local network reliability and performance are improved by regular enhancements to the technology, architecture, and integrated network management and monitoring services. Firewalls and an intrusion detection/prevention system co-managed with the Computer Security Program form the first line of defense for information security. Remote-access services enable reliable, secure access for staff and collaborators to the Laboratory's computing resources. Laboratory staff benefit from a computing ecosystem that allows them to perform their technical and business activities successfully from the desktop.

Centralized Services and Accomplishments in 2003

Calendaring, email, ph (e.g., white pages), and Entrust encryption are among the centralized services provided. To protect computing resources, internal and external (i.e., to/from Internet), email is scanned for viruses, and external email is scanned for spam. Users require fewer passwords with centralized identification management for authentication now used by many business services at the Laboratory. To ease legitimate file sharing and protect information, development efforts are underway to demonstrate role management by enforcing a common set of business rules for granting access to sensitive information.

Calendar 2003 was a productive year for the staff developing and supporting the numerous centralized network services. Specific and significant accomplishments included the following.

- Migration to the new backbone was completed, Internet access was upgraded to 622Mbps, and the path between both was upgraded with redundant dual Gigabit Ethernet firewalls.
- Email was enhanced with improved metrics gathering and an anti-spam service.
- The server and clients for the centralized calendaring service were upgraded.
- The email list management service was upgraded to an enhanced new product.
- An exceptional multi-directorate effort resulted in just under 100 business applications using the centralized username and password database.
- Remote access services ("Best in Class" in the Gartner Survey) were enhanced to use One-Time Passwords (OTP).

Growth in LLNL Network Services, 1999–2003

	1999	2000	2001	2002	2003
Percent uptime for backbone	99.87%	99.77%	99.90%	99.84%	99.94%
Number of backbone connections	161	170	181	209	219
Number of network attached devices	25,720	27,551	33,653	36,854	46,420
Approx. weekly data forwarded by backbone	NA	NA	NA	NA	19.1 TB
Approx. weekly data to/from Internet	NA	NA	NA	NA	2.6TB in / 1.8TB out
Percent recevied email marked as spam	NA	NA	NA	NA	23%
Email virus scanner					
External messages scanned / viruses detected	NA	NA	10,878,555 / 31,040	16,401,745 / 119,170	18,825,825 / 829,498
Internal messages scanned / viruses detected	NA	NA	23,937,072 / 175	56,287,228 / 687	59,487,334 / 5,915
Number of central email users	9,694	9,749	10,972	12,531	12,936
Number of central POP users	8,033	8,156	8,809	8,913	10,080
Number of central calendaring users	NA	5,609	7,893	8,728	9,149
Number of Entrust (encryption) users	NA	1,684	2,750	2,940	3,378
Remote access for staff					
Number of VPN staff accounts	NA	NA	NA	NA	3,176
Number of VPN logins by staff	NA	NA	61,266	123,666	135,360
Number of IPA logins by staff	541	75,030	137,709	101,767	54,709
Number of other remote access (e.g., modem) logins by staff	NA	NA	NA	47,856	115,762
Remote access for collaborators					
Number of collaborator remote access accounts	NA	NA	NA	NA	726
Number of VPN&VPN-C remote access logins by collaborators	NA	NA	NA	182	5,737
Number of IPA remote access logins by collaborators	NA	NA	NA	6,008	20,323

Figure 2.06-1. Growth of network services has been dramatic for several years and continued in FY03.

Problem Description

The LLNL Computer Security Program (CSP) was established in February 2000. During 2003, CSP made significant progress toward enhancing the security of LLNL's computing resources. Every day, the Laboratory is bombarded with attempts to attack its internal networks and computers from the Internet. At the same time, there is an explicit DOE requirement to maintain the security of the Laboratory's infrastructure from possible insider threats. With the advances in computing technology and the growth in the number of Internet users, successfully solving these problems is difficult and ongoing, especially since the methods used to attempt unauthorized access are rapidly changing.

Technical Approach/Status

The CSP employs a variety of approaches to coordinate and manage cyber security functions, and to protect LLNL networks and computer infrastructure from intrusion and attack. Network protection is enhanced through the use of an Intrusion Detection and Response (IDR) fabric that overlays all three networks. Through risk analysis, assessments, and management, risks are identified and an optimum risk mitigation strategy is determined. At the same time, CSP employs an active vulnerability scanning function across all three networks as well as doing wireless "war driving" and modem "war dialing." Users with detected vulnerabilities are instructed to fix or mitigate the vulnerability, or have their network access disconnected. In addition, the CSP provides technical expertise to Laboratory programs to help solve and resolve, as appropriate, computer security problems or issues that are impacting their operations.

Progress in 2003

Major upgrades of the IDR systems on the Yellow network now accommodate OC-12 and Gigabit Ethernet speeds. Significant intrusions into our computer networks were prevented thanks to vulnerability scanning and identifying problems before they could be exploited.

We completed a risk assessment of voice-over-IP technologies to enable deployment on LLNL's classified network. After conducting the risk assessment, we received approval from DOE to remove the requirement for Personnel Security Assurance Program (PSAP) access authorization for File Interchange Systems (FIS) access.

Significance

The continued progress made by the CSP is best characterized by the fact that there has not been a known intrusion into a computer on our network in more than 12 months. Through the IDR fabric and the virus and malicious code detection capabilities that are present on the networks, the Lab has not experienced a major disruption from malicious code and LLNL's users rarely have to worry about any intrusions. This "defense-in-depth" security infrastructure allows LLNL employees to concentrate on their work assignments instead of computer security.

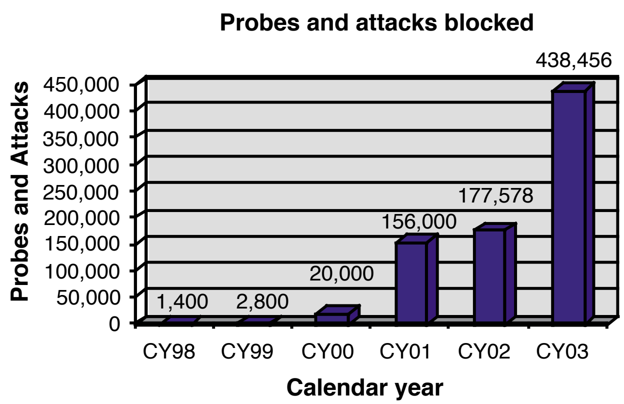

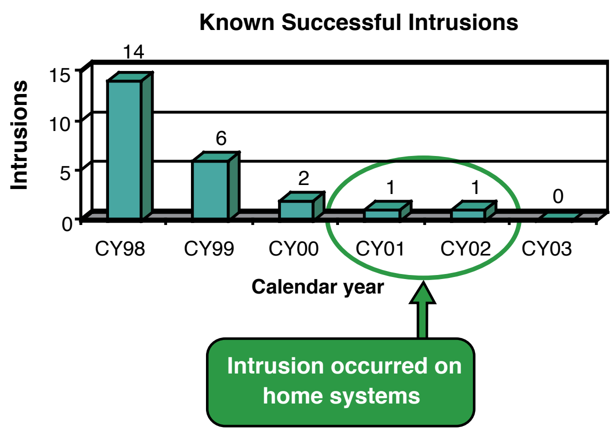

Figures 2.07-1 (Probes) and 2.07-2.(Intrusions). Large-scale coordinated attacks were detected but were unsuccessful. No known successful intrusions occurred in 2003. Despite the high number of attempted intrusions, LLNL's network continues to be secure.